I'll Be Watching You:

Intelligent Vision Surveillance Systems in the United States of America

Mila Grandes, Estefanía Szprengiel, Mitch Lenzen, and Asim Varma

School of Continuing Studies, Georgetown University

MPIM-5000-102: Ethical AI

Dr. Pablo Molina

April 19, 2026

Abstract

This paper describes computer vision surveillance in the United States and its evolution of Intelligent Vision Surveillance Systems (IVSS): integrated, scalable, and inferential systems that make people operationally legible across time, space, and databases. Drawing on commercially available tools such as facial recognition, image retrieval, behavioral analysis, and data-fusion platforms, the paper shows how surveillance is shifting from passive observation to unified architectures capable of identification, retrospective search, tracking, and institutional action. It examines the harms produced by this shift, including erosion of anonymity, false identification, discriminatory bias, cybersecurity vulnerability, and chilling effects on speech, protest, and association. The paper further evaluates these systems through utilitarian, deontological, and moral foundations frameworks, arguing that public safety claims alone cannot justify the burdens IVSS impose on privacy, autonomy, and democratic participation. Finally, it compares European and U.S. regulatory approaches and contends that the current American legal framework remains fragmented and inadequate for systems built for persistent, integrated surveillance. The paper concludes that the traditional privacy-versus-security frame is outdated and that stronger, rights-protective regulation is necessary.

Intelligent Vision Surveillance Systems in the United States of America

The most consequential development in contemporary surveillance is not any single breakthrough in computer vision, but the growing ability to assemble commercially available tools into unified systems that can monitor, identify, retrieve, and act on people across time, space, and databases. As products marketed for facial recognition, image search, behavioral

analysis, and data fusion increasingly offer complementary capabilities, surveillance shifts from passive observation toward an integrated and scalable architecture.

In turn, that architecture produces compounding harms: it erodes anonymity, increases the risk of false identification, reproduces racialized and discriminatory patterns of surveillance, creates cybersecurity vulnerabilities, and chills speech, protest, and association by expanding institutional power. At the same time, examining mass surveillance through utilitarian, deontological, and moral foundations frameworks shows that public safety claims alone cannot resolve the deeper conflict between collective security and fundamental rights. The legal problem is equally significant; there is no comprehensive federal law, and the existing U.S. constitutional framework was largely developed for more discrete and limited surveillance and has not kept pace with systems built for persistent monitoring, retrospective search, and cross-platform data integration.

How IVSS Works and How It is Different

Defining Surveillance and Computer Vision Surveillance

Human surveillance is the gathering of information about a person through observation in order to monitor their presence, identity, behavior, relationships, and movements. Computer vision is the branch of artificial intelligence concerned with visual data. Kalluri et al. (2025) define it as "AI that focuses on measuring, recording, representing, and analyzing the world from visual inputs such as image and video data" (p. 73). Thus, computer vision surveillance is not merely the observation of human behavior; rather, it is a process through which artificial intelligence transforms human presence into machine-readable, operationally useful information.

Linguistic Obfuscation for Terminology Disorientation

Scholars have observed that academic papers and patents in this area generically refer to images, text, or objects, leaving unstated the anticipated use of computer vision technology for surveillance by the government (Kalluri et al., 2025). They further note that "only 1% of papers and 1% of patents were dedicated to extracting only non-human data," a finding that reveals how extensively computer-vision research and its applications are involved in "datafying humans, specifically human bodies" (Kalluri et al., 2025, p. 75). This is not a minor rhetorical issue. These findings suggest that surveillance-oriented research is not marginal or accidental, but pervasive, normalized, and often obscured by generic or depersonalized language.

The prevalence of obfuscating language in computer vision raises a further problem: if surveillance can be hidden behind neutral references to "images," "objects," or "scenes," it can also be concealed by broader and less politically or ethically charged labels for the systems themselves. What was once more directly legible as computer vision surveillance can now appear under the broader and more neutral label of Intelligent Vision Systems (IVS), a term that reflects the field's contemporary emphasis on automated, adaptive, and data-driven visual technologies.

How IVSS Differs From Computer Vision Surveillance

The long-running Advanced Concepts for Intelligent Vision Systems (ACIVS) conference describes its field as concerned with the development of "adaptive, intelligent, safe, and secure imaging systems" (ACIVS 2002, 2002). In the absence of a single settled academic definition of Intelligent Vision Systems (IVS), the ACIVS conference record provides a useful basis for defining the term. Across that record, IVS appears less as a fixed category than as an expanding set of technical capacities. Early work centered on image processing, recognition, retrieval, and system performance; later research extended into multi-camera systems, smart

environments, human-centered analysis, biometrics, forensics, and deep-learning-based interpretation (see Appendix A). On that basis, IVS can be defined as systems that capture, process, connect, and interpret audiovisual data in increasingly adaptive and intelligent ways. IVSS, in turn, can be defined as the surveillance-focused development of those capabilities into increasingly integrated and pervasive systems for monitoring, analyzing, and retrieving people-related information.

In this sense, IVSS can be understood as the technical maturation of computer vision surveillance. However, these systems differ from prior surveillance technologies because they operate at unprecedented scale, inference capability, and integration with identity systems. IVSS is designed to detect, classify, track, identify, retrieve, and analyze people across time, space, and databases, whether through CCTV, smart-camera networks, biometric extraction, person search, or retrospective retrieval architectures.

Commercial Off-The-Shelf Products

Within a short period of time, private companies have developed a number of products with increasing sophistication and features that they market to law enforcement.

Corsight and Fortify

Corsight AI markets Fortify as a facial-recognition and contextual visual-analysis platform for law enforcement, public safety, and mass-gathering venues (Corsight, n.d.-a) Corsight (n.d.-c) describes the system as supporting the identification of wanted criminals and missing persons, large-scale post-event investigation, advanced image matching across local and national datasets, and real-time deployment from control rooms, vehicles, tablets, and mobile devices., Corsight (n.d.-a) further presents the platform as being designed to operate in crowded

and difficult environments, including low-quality video, extreme angles, motion, darkness, and partially covered faces, while integrating with existing camera infrastructure (Corsight, 2025b). These materials present Fortify as a deployable surveillance product capable of real-time watchlist matching, retrospective video review, and large-scale database search as Fortify can connect to any database (Corsight, n.d.-b).

Figure 1

Corsight AI's Identification of a suspect after an incident

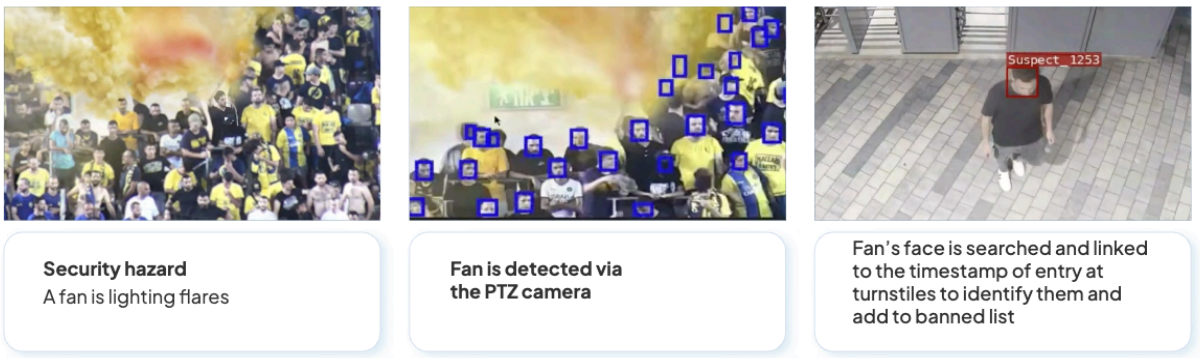

Note: Subject identification using a Pan-Tilt-Zoom (PTZ) camera (Corsight, 2025a).

Clearview and its 70 Billion Image Database

Clearview AI markets a web-based facial-recognition platform built around a very large facial-image repository and designed primarily for law-enforcement investigation (Clearview AI, n.d.-a). Clearview's platform allows agencies to upload a facial image to compare it to the more than 70 billion public online images and any custom galleries selected by the agency to identify matches and source URLs in seconds (Clearview AI, n.d.-b, n.d.-c). Clearview states that its facial image repository consists of images drawn from social media posts, personal and professional websites, news articles, mugshot sites, public records sites, and other sources. Clearview markets the platform and frames the system as a scalable tool for generating

investigative leads, identifying persons who might otherwise remain unknown, and supporting watchlists, alerts, and identity management (Clearview AI, n.d.-c, n.d.-d).

Figure 2

Search Workflow to Retrieve an Identification Match from Clearview

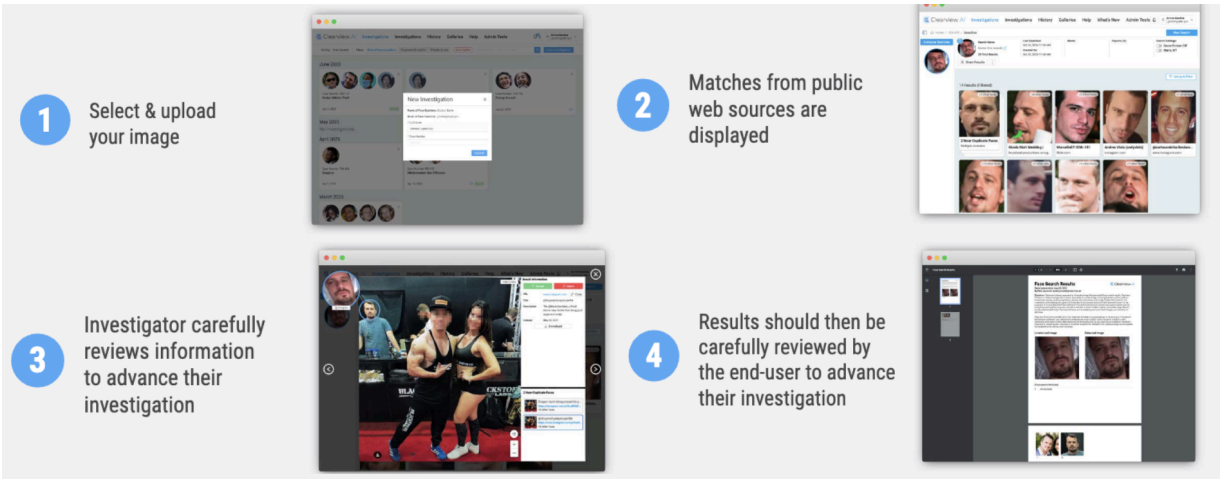

Note: The steps above detail the process to obtain an identity match from submitting an image to Clearview's database. From one query, nine images are displayed, portraying a 1:9 ratio of person to collected and retrieved images in Clearview's database (Clearview AI, n.d.-d).

Palantir

Palantir markets Gotham as a centralized system for bringing together data from disparate sources so authorized users can search, analyze, and share it within a single controlled environment (Palantir Technologies UK, Limited, 2024). Palantir also states that Gotham can process multiple forms of data, including audio and video, and support large-scale investigative analysis such as link analysis, geospatial analysis, object analysis, call-detail-record analysis, and alerting (Palantir Technologies UK, Limited, 2024).

Gotham's Browser and Custom Object Views tools allow users to search across integrated and federated data, inspect entity-level records and relationships, configure dashboards for specific workflows, and receive alerts when search criteria are met or tracked objects change state or location (Palantir Technologies UK, Limited, 2024). On its public-facing pages, Palantir describes this broader architecture as "AI-Powered Automation for Every Decision" and, through its Ontology system, as a platform that encodes the data, logic, action, and security of the enterprise in order to automate decisions across operations (Palantir Technologies Inc., n.d.-a, n.d.-b).

Figure 3

Palantir's Custom Object Views

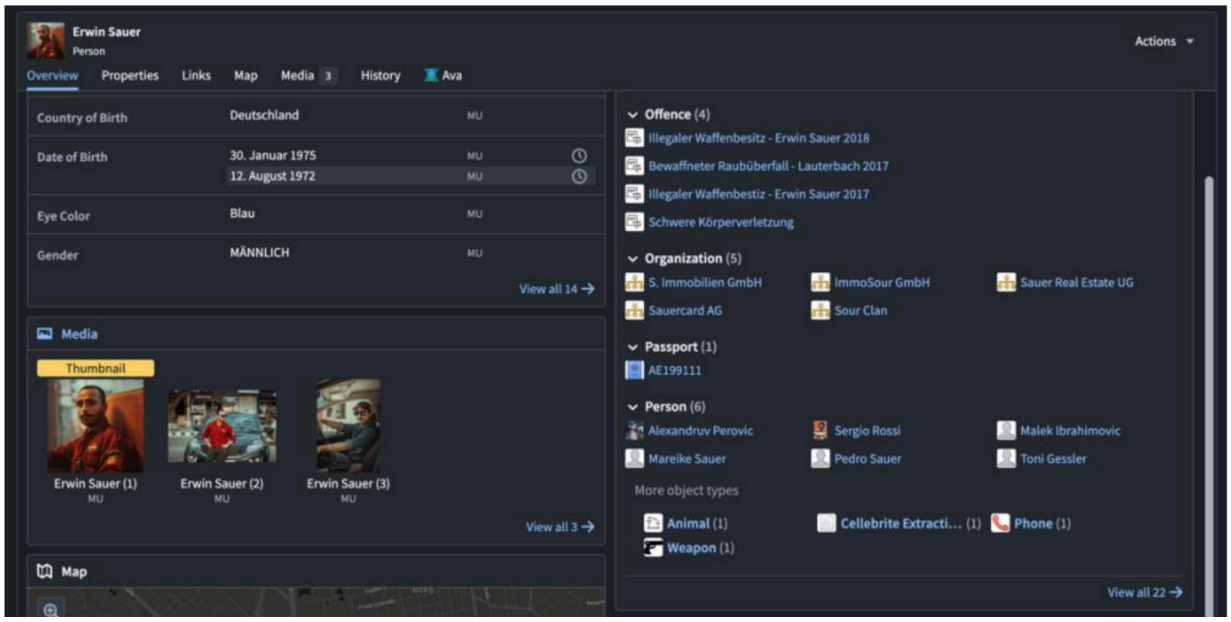

Note: The image above displays the profile of an individual, which includes multiple records that encompass the known information about the person.

Palantir is significant because it provides the integrative layer through which many surveillance functions can be fused, queried, governed, and operationalized. Its materials

describe a system capable of unifying multimodal and real-time data, creating a single source of truth, linking records about persons and objects, generating search and geofence alerts, and supporting intelligence production, investigative network analysis, mission planning, execution, and after-action review (Palantir Technologies UK, Limited, 2024).

The government was seeking just such a system. A 2025 ICE procurement rationale stated that it required a modified COTS solution capable of combining investigative case management, data analytics, an enterprise lakehouse repository, interoperability with internal and external law-enforcement and biometric-related systems, media management, role- and case-level security, mobile and offline access, and that its market research suggested Palantir was the only vendor able to meet those requirements within the necessary timeline (U.S. Immigration and Customs Enforcement, 2025).

Building a Unified Mass Surveillance Pipeline

Corsight, Clearview, and Palantir illustrate how IVSS can be assembled into a unified, end-to-end surveillance pipeline from commercially available products. In such a configuration, Corsight would function as the front-end perceptual layer, continuously monitoring camera feeds for faces, persons of interest, and suspicious patterns of behavior; Clearview would function as the identity-enrichment layer, taking captured face images and comparing them against a vast searchable archive of public-web and custom-gallery images; and Palantir would function as the ontological and operational layer, integrating these outputs with case records, geospatial information, media, alerts, relationships, and institutional workflows into a common operating picture (Corsight, n.d-a.; Clearview AI n.d.-a, n.d.-b; Palantir Technologies UK, Limited, 2024; Palantir Technologies Inc., n.d.-b). This is an analytical synthesis rather than a claim about one single confirmed deployment.

The capabilities marketed across these systems are technically complementary and if combined would form a coherent architecture for mass surveillance. Such a pipeline moves surveillance beyond passive observation and toward near-total operational legibility. It would reduce anonymity in crowds, merge real-time and retrospective surveillance, and connect sensor-based monitoring to large-scale data fusion and institutional action. In practical terms, it could support suspect identification, missing-person location, ban-list enforcement, behavioral flagging, retrospective archive search, geofenced tracking, and network analysis within one interoperable framework. An all-seeing surveillance pipeline would not require a technological breakthrough. It could be built from already available commercial layered products.

Figure 4

A Hypothetical Unified Mass Surveillance Pipeline

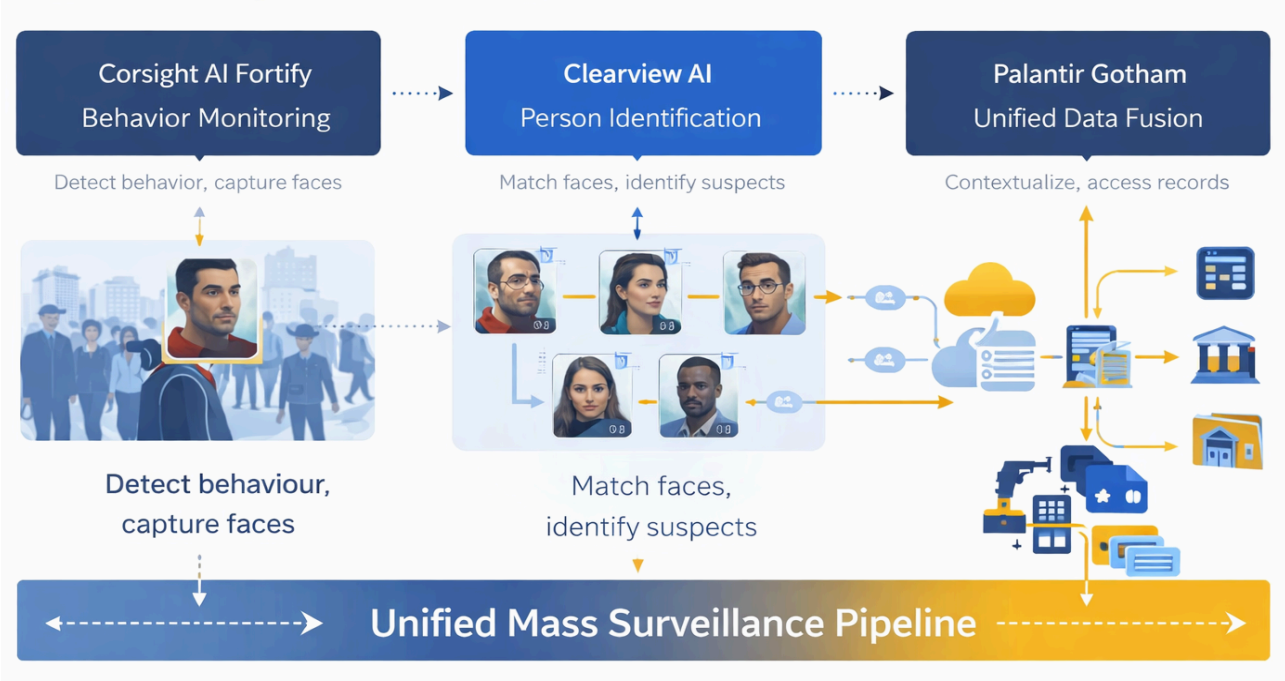

Note: This a conceptual model illustrating how commercially available systems could be combined into an end-to-end mass-surveillance architecture.

Harms: The Security Paradox in Practice

Intelligent Vision Surveillance Systems (IVSS), encompassing facial recognition technology (FRT), automated license plate readers, predictive analytics platforms, and integrated data brokerage tools, have been deployed rapidly and with minimal regulation by law enforcement agencies across the United States and beyond. Their proponents frame them as indispensable public safety infrastructure. Yet this security rationale conceals a paradox: the very systems deployed in the name of public safety serve as the infrastructure that has the potential to harm the public by eroding privacy, producing inaccurate identifications, amplifying existing biases, and creating cybersecurity vulnerabilities. In the hands of an authoritarian government these systems can suppress the civic freedoms democratic governance is meant to protect.

The magnitude of this issue encompasses a large swath of the U.S. population, many of whom have no idea they are enrolled in what Garvie, Bedoya, and Frankle called "The Perpetual Line-up" (Garvie et al., 2016). These profiles, drawn from state ID registries, federal travel documents, and criminal records, mean that roughly half of American adults have had their photos enrolled in a facial recognition network searchable by law enforcement (Garvie et al., 2016). At least 26 and possibly up to 30 states allow law enforcement to search these repositories, which are primarily DMV registries built from photos submitted by law-abiding citizens for identification rather than forensic purposes (Garvie et al., 2016). Beyond DMV records, these databases also include mugshots and are not limited to convicted criminals; they frequently contain images of people whose charges were dropped, who were acquitted, or who were arrested for nothing more than a misdemeanor, such as peaceful civil disobedience (Garvie et al., 2016). This shift marks a radical departure from historical practice, as biometric databases once consisted almost exclusively of information from criminal arrests or forensic investigations

(Garvie et al., 2016). In this perpetual line-up, an algorithm, rather than an eyewitness, identifies suspects among a population that is overwhelmingly non-criminal.

This section analyzes five aspects of real-world harm and their chilling effects on civic engagement. Peer-reviewed research (Garvie et al., 2016; Gentzel, 2021; Marx, 2016; Murray et al., 2024), policy research (Turner Lee & Chin-Rothmann, 2022), and verified investigative journalism (Franceschi-Bicchierai & Whittaker, 2026; Frenkel & Krolik, 2026; Business & Human Rights Resource Centre, 2026), collectively demonstrate that these are mutually reinforcing features of IVSS that threaten privacy rights, democratic participation, and the rule of law.

Privacy Erosion

IVSS shifts surveillance from the passive observation of human activity to the mass, automated, real-time monitoring of all activity. Traditional cameras are transformed into active "brains" that can identify and evaluate human behavior frame by frame. By fusing an individual's physical presence with their digital identity, IVSS effectively dissolves the "practical obscurity" that has characterized movement through public spaces (Stanley, 2019). As a result, the traditional expectation that our daily actions will not be scrutinized is replaced by a condition in which public behavior is subject to constant evaluation and judgment by agents of authority.

This erosion of anonymity poses a profound threat to democracy's core pillars, since privacy is an essential "expression" of a person's engagement in society and is necessary for human flourishing and the maintenance of trust (Tavani, 2016). Anonymity in public spaces provides the psychological safety required for free expression and dissent. However, as individuals become aware that their actions are being scrutinized and recorded, IVSS can chill that engagement. In the United States, law enforcement agencies are already using IVSS features

to identify undocumented immigrants. For example, Immigration and Customs Enforcement (ICE) uses a "tech arsenal" that includes Clearview AI and Palantir's ELITE tool to map neighborhoods, generate detailed dossiers, and assign "address confidence scores" to guide raids (Franceschi-Bicchierai & Whittaker, 2026; Cox, 2026). ICE agents in Minneapolis have also reportedly used Corsight's Mobile Fortify and other facial recognition tools to identify and track citizens protesting the agency's presence, often without meaningful legal oversight or judicial warrants (Frenkel & Krolik, 2026).

For ordinary, law-abiding people, the implications of this toolkit are significant: every photograph or opinion posted on social media, every check-in at a restaurant or march, and every "like" given to a political cause can be assembled into a comprehensive dossier of a person's beliefs, associations, relationships, and movements. What was offered to the world as self-expression, community building, or civic participation can, in the hands of a sufficiently motivated authority, be reprocessed as evidence of identity, ideology, or dissent. The individual who posted a photo at a protest, tagged their location at a place of worship, or publicly followed an advocacy organization did so under the assumption that expression and surveillance were separate domains. As Tavani (2016) warns, once personal information enters the digital ecosystem, the individual loses meaningful control over its future use; the data persists, accumulates, and becomes available to institutions whose interests may be fundamentally opposed to those of the person who generated it.

Ultimately, the loss of privacy enabled by IVSS is likely irreversible, since privacy, once lost, is nearly impossible to reclaim (Tavani, 2016). In the digital age, these systems contribute to what has been termed a "womb-to-tomb dossier": a permanent electronic record that follows an individual throughout life and makes it impossible to "start over with a clean slate" (Tavani,

2016). This reality aligns with the structural logic of the Perpetual Line-up identified by Garvie et al. (2016). Just as more than 117 million Americans are enrolled in facial recognition networks without their knowledge or consent, the aggregation of social media data, location signals, and biometric identifiers means that virtually every digitally active adult is similarly drawn into a surveillance infrastructure they never agreed to join. The face scanned by Corsight's Mobile Fortify on a Minneapolis street, the phone intercepted by a cell-site simulator, and the social media profile compiled by Penlink are not separate phenomena; they are converging streams feeding a single dossier.

The case of Chris Wilson, a Black activist and scholar, illustrates how modern biometric infrastructure can facilitate the permanent and systemic erosion of privacy, turning a single encounter with law enforcement into lifelong harm. After Wilson was arrested for a minor misdemeanor during a peaceful protest in 2016, her mugshot was enrolled in both the Federal Bureau of Investigation's Next Generation Identification database and the Pinellas County Sheriff's Office facial recognition system (Garvie et al., 2016). Now, when Wilson participates in a protest, her information appears to law enforcement as they surveil the event, and she immediately becomes a "person of interest" (Garvie et al., 2016). Her story also underscores the permanence of digital records: once an individual is categorized as a "subject of interest," they are trapped in a perpetual state of scrutiny that chills their free expression in public.

As these surveillance infrastructures become embedded in the urban fabric, personal identity is seized as a form of digital capital, often through the non-consensual capture of biometric "face prints." Forcing individuals to exist under perpetual scrutiny chills democratic participation and facilitates a form of techno-panoptic governance that drives de-socialization (Murray et al., 2024; Oztig & Karluk, 2025). This illustrates a security paradox: tools intended to

promote safety can fundamentally erode privacy while undermining civil liberties. Social media platforms were similarly designed as spaces for connection, self-expression, and community building. Yet within the surveillance infrastructure described above, it is worth considering what individuals voluntarily surrender each time they post a photograph, share their location, declare a political opinion, or publicly affiliate with a cause. Each act of digital participation can hand law enforcement a detailed map of a person's network, revealing not only that person's beliefs and movements but also those of their friends, family members, and associates. If that awareness is already causing people to think twice before posting, attending a public event, or openly associating with a cause, then the chilling effect is no longer a future concern. It is a present reality.

Inaccuracies and Discriminatory Biases

The inherent inaccuracies of computer vision identification and law enforcement's overreliance on automated outputs transform IVSS into a tool of active harm. The human cost of technical fallibility is most poignantly illustrated when algorithmic errors are amplified by systemic failures in police procedure, transforming investigative leads into life-altering catastrophes.

The case of Angela Lipps, a Tennessee grandmother who spent over five months in jail for crimes in a state she had never visited, underscores the devastation of using technology as a shortcut for basic investigation (Sottile, 2026). Despite exculpatory bank records being readily available, a faulty Clearview AI match led to Lipps being terrified, exhausted, and humiliated while wrongfully extradited and imprisoned halfway across the country (Sottile, 2026). Similarly, Taki Allen, a Maryland high school student, was handcuffed at gunpoint after an AI security system mistook an empty bag of chips for a firearm. Although human reviewers

correctly identified the error and canceled the alert, a breakdown in communication between school officials and law enforcement resulted in a traumatizing tactical response involving eight police cars (Tsui & Sottile, 2025). These two cases exemplify the physical and psychological danger inherent in automated suspicion. In both cases, the wrongfully accused parties reported long-lasting negative emotional and, in the case of Lipps, also financial effects (Sottile, 2026).

It is essential to consider whether the speed at which law enforcement is deploying AI surveillance has outpaced the human capacity for critical oversight, effectively turning citizens into collateral damage in an unproven experiment of algorithmic policing. As researchers have noted, agencies often rely on vendor promises with little evidence of efficacy, leading to nightmare scenarios when law enforcement agents seek shortcuts in algorithmic results (Sottile, 2026). While AI-driven identification and detection tools are not inherently malevolent, the training of both the systems and the personnel overseeing them must be rigorously refined before deployment. Society must grapple with the reality that without such safeguards, these tools will continue to inflict lasting psychological harm and the unjust loss of freedom on law-abiding citizens.

Furthermore, a landmark research by Buolamwini and Gebru (2018) revealed that commercial algorithms exhibit staggering error rates: while light-skinned males experience a near-negligible 0.8% error rate, misclassification for dark-skinned females reaches as high as 34.7%. Overall, these systems perform 11.8% to 19.2% worse on darker-skinned subjects, a disparity often attributed to the "training is destiny" phenomenon, where datasets predominantly feature Caucasian faces (Gentzel, 2021; Garvie et al., 2016). Exacerbating the disparity in accuracy is the bias observed in facial recognition tools that use emotion analysis, which tends to interpret Black faces as expressing ”anger" or "contempt" at higher rates than white faces

(Gentzel, 2021). IVSS, using biased algorithms and interpretations of malintent, is used to justify “preemptive law enforcement” even though it is completely without scientific basis. The use of such biased technology by government entities is inconsistent with the requirement that all citizens be treated equally before the law (Gentzel, 2021). By failing to account for demographic variance and the complexities of real-world deployment, these systems may facilitate wrongful identifications that carry life-altering legal consequences.

Real-world conditions further degrade system accuracy. An experiment at a train station in Mainz, Germany, revealed that while identification accuracy may hold at 60% during daylight, it plummets to 10–20% in darker lighting conditions (Garvie et al., 2016). Companies like Clearview AI attempt to mitigate this by utilizing "mathematical maps" of facial traits to claim high accuracy even in poor light (Hill, 2024). However, such claims often foster a dangerous "automation bias" among law enforcement officers. As Kashmir Hill (2024) observes, an over-reliance on technology can lead officers to believe they can solve crimes "without leaving their desks," neglecting traditional interrogative methods that might otherwise disprove a false match (p.235). Given the rate of inaccuracy, results from facial recognition technology (FRT) searches should not be the only basis for an arrest, and reportedly, they are not, however as the following case suggests, FRT still plays a persuasive role in arrests.

The case of Robert Williams, the first man in the U.S. to be wrongfully arrested due to a racially biased algorithm, serves as the definitive warning. Williams was detained for 30 hours based on a faulty match from security footage and auxiliary information regarding a previously pawned watch. Viewed through the biased lens of the algorithmic match, detectives interpreted this mundane detail as incriminating evidence, demonstrating how these 'investigative leads' often serve as the primary framing for an arrest despite official policies to the contrary (Allyn,

2020; Hill, 2024). The limitations of IVSS are not merely engineering "glitches"; this biometric infrastructure creates a systemic burden that falls disproportionately on people of color due to a combination of technical inaccuracies, predictive policing, and over-indexing in criminal databases.

Communities of color, often concentrated in densely monitored urban or low-income areas, are subjected to a higher volume of "invisible" scans than more “affluent and whiter” areas. This is the result of predictive policing programs, where algorithms make a geographical prediction of where crime is most likely to happen based on historical crime data (O’Neil 2016). Because low-income, urban communities of color are subjected to a higher density of cameras, they are over-indexed in surveillance databases regardless of criminal history. In the context of IVSS, this means that the populations for whom the technology is least accurate are also the most frequently scanned, and consequently, this demographic becomes overrepresented in the database. The "perpetual line-up" is not a race-neutral occurrence, and it is built upon historical and contemporary patterns of racialized surveillance.

Security and Cybersecurity Risks

The centralization of biometric, behavioral, and personal data, now potentially extending to DNA samples (Anderson, 2026), transforms IVSS databases into high-value targets whose compromise would expose entire populations to harms of unprecedented scale. The risks are twofold: internal misuse by the agencies entrusted with this data, and external exploitation by malicious actors. In both cases, the public bears the consequences of governance failures it has little power to prevent or even detect.

Centralized surveillance databases are only as secure as the personnel who operate them. Whistleblower allegations concerning the Department of Government Efficiency illustrate this

vulnerability with alarming clarity: a former software engineer reportedly exfiltrated sensitive records on over 500 million Americans from the Social Security Administration onto a personal thumb drive, allegedly boasting of having 'God-level' access to restricted databases containing Social Security numbers, birth dates, and ethnicity data (Franceschi-Bicchierai, 2026; Larson, 2026). This is a predictable consequence of granting unfettered access to high-value data without standardized auditing or meaningful accountability. Garvie et al. (2016) found that in law enforcement operating IVSS systems, fewer than 10% maintain publicly available use policies; the question of whether personnel will follow protocols, rather than exploit access for personal gain, becomes a structural governance failure.

Beyond the risks posed by internal actors, the structural vulnerability of centralized data systems is being exacerbated by a new generation of AI-driven cyber threats. The rise of sophisticated "agentic" AI models has dramatically lowered the barrier for breaching complex security infrastructures, as these tools can reason, improvise, and execute tactical operations autonomously at an infinite scale (VandeHei, 2026). Tech leaders have privately warned that upcoming models, such as Anthropic's "Mythos," may be "scary good" at hacking sophisticated systems, presaging a wave of attacks that far outpace the current defensive capabilities of even the most well-funded corporations and municipal governments (VandeHei, 2026). The prevailing strategy of consolidating vast quantities of citizens' sensitive information into centralized, virtual databases—such as the SSA's master database or ICE's ELITE system—creates extraordinarily high-value targets for malicious state-sponsored actors and cybercriminals (Fowler & Joffe-Block, 2026; Cox, 2026). As a security risk to an AI system inevitably creates safety lapses that can harm humans, the commodification of this data into

behavioral dossiers ensures that any breach represents a structural collapse of the national identity and security system (Steitz, 2026; Larson, 2026).

The same features that make IVSS powerful from a law enforcement perspective, like their aggregation of vast biometric and behavioral datasets, their integration of government and commercial data streams, and their reliance on third-party vendor software and hardware, simultaneously make them high-value targets for cyberattack. One of the most acute and vulnerable points of entry is the very technology at the heart of computer vision, the surveillance camera. In this broader data ecosystem is the proliferation of surveillance camera networks. Unauthorized access to the feeds is no longer limited to viewing video footage; with the advent of "edge computing," the feed now transmits digital representations of faces, vehicle data from license plate readers, and detailed behavioral analysis directly to intruders (Béchard, 2026). An extreme case of the exploitation of surveillance camera vulnerabilities is seen in conflict zones where hacked surveillance feeds have been used to guide targeted assassinations and track the movement of convoys (Béchard, 2026). Part of the vulnerability, as Béchard (2026) reports, comes from the difficulty of routinely updating the software and the physical cameras, which are oftentimes left for years exposed to a myriad of new computer viruses and hacking techniques before getting a software upgrade.

When technology providers explicitly disclaim responsibility for the performance and accuracy of their products, law enforcement agencies are stripped of the contractual leverage necessary to mandate rigorous security audits, vulnerability disclosures, or essential software updates (Garvie et al., 2016). This absence of accountability is a pervasive feature of the modern surveillance ecosystem, where agencies frequently adopt powerful artificial intelligence tools based solely on vendor promises rather than documented evidence of efficacy (Sottile, 2026). It

is critical to recognize two divergent accountability gaps created by these contractual shields. In instances of wrongful arrest, such as the incarceration of innocent individuals due to misidentification, private vendors typically are not liable because the specific authority to deprive a citizen of liberty rests exclusively with the state (National Academies, 2024). Consequently, the responsibility for unlawful detentions remains with the police department, regardless of the faulty algorithmic lead that initiated the investigation (Garvie et al., 2016; Gentzel, 2021).

In the realm of cybersecurity, the structural gap is even more profound because a single vendor may hold biometric data for millions of law-abiding citizens while maintaining a disclaimer of all warranties regarding the reliability of their system (Garvie et al., 2016). This creates a scenario where the public bears the entirety of the risk while the vendor remains insulated from the consequences of potential breaches or exploits (Garvie et al., 2016). As experts have warned, once sensitive personal data is leaked due to such systemic failures, it constitutes an irrecoverable loss with generational consequences for the national identity system (Larson, 2026). These failures raise fundamental questions for society: should the state be permitted to outsource the mechanics of justice to private firms that refuse to stand behind the reliability of their own tools (Gentzel, 2021)? Furthermore, given the known inaccuracies of IVSS, what steps must law enforcement authorities grapple with whether their reliance on these disclaimed systems constitutes a breach of the public trust, as it effectively transforms the citizenry into a permanent and non-consensual testing ground for unproven and unbonded technology (Gentzel, 2021; Sottile, 2026).

A breach in the trust is the use of cell-site simulators, which are devices that mimic cell towers to force nearby phones to connect, thus they enable identification of the owner,

interception of their calls, messages, and location data (Franceschi-Bicchierai & Whittaker, 2026). These add another layer of supply-chain vulnerability to the IVSS ecosystem. ICE has signed contracts exceeding 1.5 million dollars with a company that integrates cell-site simulators into specialized vehicles for law enforcement use (Franceschi-Bicchierai & Whittaker, 2026). These devices are designed to capture the data of all nearby phones, not only those belonging to targets, and have been deployed without warrants in documented cases. More troublingly, prosecutors have actively concealed the use of cell-site simulators to extract plea deals or drop cases rather than disclose the technology or the non-disclosure agreements governing it.

The convergence of these internal and external threats underscores the urgent necessity for robust data governance, comprehensive transparency mechanisms, and enforceable security frameworks. However, the current administration’s posture appears to be prioritizing the removal of safeguards over the construction of the necessary security infrastructure to mitigate these risks. This is most evident in the ongoing dispute with AI vendors over usage restrictions, where the Department of War has threatened to designate companies as "supply chain risks" if they refuse to remove guardrails against mass domestic surveillance and fully autonomous weapons (Amodei, 2026). By engaging in "fishing expeditions" for fraud and demanding unprecedented, unfettered access to sensitive records without judicial oversight, the administration is effectively accelerating the deregulation of a technology that is not yet reliable enough to protect the fundamental rights and safety of the American public (Amodei, 2026; Garvie et al., 2016). This failure to establish accountable deployment models signals a move toward a "creeping authoritarianism" where citizen data is treated as a disposable resource rather than a sacred trust, a theme that will be analyzed in greater depth in the following sections regarding legislative oversight and policy reform.

Chilling Effects on Civic Engagement

Surveillance patterns have historically targeted those who challenge the status quo, from the FBI's tracking of civil rights leaders like Martin Luther King, Jr., and Malcolm X to the contemporary misuse of facial recognition technology (FRT) against activists. Modern IVSS enables a more "precise discrimination" by allowing law enforcement to build harmful cycles of monitoring that disproportionately impact those who have challenged authority.

The harms analyzed in preceding sections, such as privacy erosion, false identification, discriminatory burden, and cybersecurity risk, are, in a meaningful sense, proximate harms: they fall on identifiable individuals in documentable ways. The chilling effect represents a different, and in some respects more insidious, category of harm: a diffuse, structural suppression of the political and civic behaviors that democratic self-governance requires. When people modify what they say, where they go, whom they associate with, and whether they participate in public life because of the awareness, or even the suspicion, of continuous surveillance, the harm is borne not only by those individuals but by the democratic community as a whole (Murray et al., 2024).

The Theoretical Foundation

Cohen (2013) warns that surveillance systems render people increasingly fixed, transparent, and predictable, a transformation that traditional privacy law, built around secrecy and individual control, was simply never designed to handle. Zuboff (2015a) sharpens this critique further: under what she terms 'surveillance capitalism,' behavioral data is not merely collected but continuously harvested and converted into prediction products engineered to forecast and ultimately shape human conduct. Together, these frameworks expose something more troubling than a privacy violation. Data-driven systems like IVSS do not passively observe behavior; they actively intervene in it. This distinction matters enormously. The gap between

what privacy law protects and what IVSS actually does is a structural problem that directly impacts democratic agency, individual autonomy, and any meaningful attempt to govern these systems (Zuboff, 2015).

Marx (2016) identifies this dynamic as a defining feature of 'new surveillance': the awareness of being watched under a system of asymmetric visibility in which the state sees the citizen but the citizen cannot see or challenge the state's gaze, and it concentrates power and suppresses political expression in ways that do not require explicit coercion. The threat of identification is itself the mechanism of control.

Murray et al. (2024), drawing on qualitative interviews with 44 activists, journalists, opposition politicians, and civil society leaders in Uganda and Zimbabwe who had been directly subjected to state surveillance, provide the most empirically grounded account available of how this suppression operates in practice. Three consistent findings emerged: first, individuals reported significant self-censorship, carefully moderating their public speech, social media posts, and associational choices to avoid drawing state attention; second, participants described an unwillingness to engage with individuals or organizations perceived as surveillance targets, out of fear of 'guilt by association'; and third, surveillance eroded the interpersonal trust essential to political organizing, fracturing movements' ability to coordinate and mobilize effectively (Murray et al., 2024).

Critically, Murray et al. (2024) found that chilling effects are not binary because individuals do not simply stop or continue engaging, but it starts producing layered modifications to behavior that aggregate into a systemic suppression of political life. The authors also note that these effects are felt most acutely at the margins of society, by those who are already in

opposition to the status quo, and that research in the United Kingdom and the United States suggests the presence of parallel chilling effects even in more democratic contexts.

This framing is legally significant. The Supreme Court established in NAACP v. Alabama (1958) that there exists a 'vital relationship between freedom to associate and privacy in one's associations,' particularly where a group advocates minority or unpopular beliefs, and reaffirmed in McIntyre v. Ohio Elections Commission (1995) that anonymity is the shield from the tyranny of the majority. Garvie et al. (2016) note that the First Amendment's protection of anonymous speech and association is fundamentally incompatible with a surveillance architecture that enables continuous, warrantless identification of any individual in public space.

Documented Suppression of Protest and Political Expression

The history of American law enforcement surveillance of political activity provides essential context. Garvie et al. (2016) document that J. Edgar Hoover's FBI conducted extensive photographic surveillance of civil rights demonstrations, deploying undercover agents disguised as freelance photographers to document protest participants, which is a practice whose explicit purpose was to identify and monitor political dissenters. In 2012, Senator Al Franken confronted the FBI about an internal PowerPoint presentation demonstrating how facial recognition could be used to identify attendees at presidential campaign rallies. In 2015, the FBI acknowledged conducting aerial surveillance over protests in Ferguson and Baltimore following police killings of Black men. The Department of Homeland Security monitored Black Lives Matter protests following Ferguson (Garvie et al., 2016). Turner Lee and Chin-Rothmann (2022) document that ICE used aerial surveillance, location tracking, and facial recognition to identify participants in the 2015 Baltimore protests, and that DHS deployed drones and helicopters over Black Lives Matter protest gatherings in at least fifteen cities following George Floyd's murder.

Contemporary Crackdowns: The Minneapolis Case

The most recent fully documented illustration of IVSS-enabled suppression of civic engagement comes from Minneapolis, Minnesota, where ICE deployed a comprehensive surveillance arsenal during its 2026 enforcement operations. Frenkel and Krolik (2026) documented, on the basis of verified videos, photographs, and testimony from three current and former Department of Homeland Security officials, that ICE used FRT not only to identify undocumented immigrants but to track citizens who participated in protests against ICE. Two photographs verified by The New York Times captured agents using Mobile Fortify to scan the faces of protesters in Minneapolis; in at least one video, agents were heard informing people that their faces were being recorded with facial recognition and that they would be added to a database (Frenkel & Krolik, 2026).

The case of Nicole Cleland illustrates how the chilling effect becomes personal. A fifty-six-year-old volunteer with a community watchdog group, Cleland was monitoring ICE activity when an agent turned and addressed her by her first name, despite the fact that the two had never met. He informed her that he had facial recognition technology and that his body camera was recording. Three days later, Cleland received an email from the Department of Homeland Security revoking her Global Entry and TSA travel privileges without explanation (Frenkel & Krolik, 2026). Cleland subsequently reflected: 'I don't know how far-reaching ICE can be. I'm struggling to figure out what I can do, without putting myself at greater risk or putting other people at risk' (Frenkel & Krolik, 2026). The government’s sanction imposed after facial recognition exemplifies the chilling sequence Murray et al. (2024) describe; surveillance becomes actionable, consequences become real, and observers modify their own behavior.

The ACLU sued DHS over the Minneapolis operations, noting that these technologies were being deployed more aggressively than before, and the convergence of all the technologies gave the government unprecedented power (Frenkel & Krolik, 2026). The danger lies not in a single tool but in the integration of FRT, Palantir's data aggregation platform, cellphone tracking, and commercial spyware into a unified monitoring architecture.

The Collective Democratic Stakes

Murray et al. (2024) frame these individual harms within a democratic analysis: surveillance erodes the trust essential for collective action. Movements fearing infiltration shrink into verified circles, limiting recruitment and mounting effective public opposition. The UN Special Rapporteur on Freedom of Assembly has affirmed that the rights to assembly and association enable individuals from groups most at risk to claim other rights and overcome challenges associated with marginalization, and that these rights must be actively facilitated rather than chilled (Murray et al., 2024). When IVSS systematically suppresses these rights it structurally disadvantages those who are already politically marginalized, concentrating power with the state at the expense of the governed.

Murray et al. (2024) argue that the current approach to surveillance analysis is fundamentally inadequate to address harms that operate at the societal level because it focuses primarily on individual privacy rights in isolation. When surveillance chills democratic participation broadly, the necessity calculus that justifies surveillance under human rights law must account for that society-wide harm, not merely for its impact on identifiable individual complainants. The evident danger, they conclude, is the emergence of creeping authoritarianism by default, in which democratic processes are hollowed out not by overt suppression, but by the systematic discouragement of the civic engagement that makes democracy real.

Freedom from surveillance is a foundational requirement for a democratic society (Cohen, 2013). Freedom from continuous identification, whether at a public march, a civil disobedience demonstration, a peaceful demonstration, a town hall, or even in the voting booth, should be an inherent right. However, navigating the "security paradox," specifically balancing the preservation of anonymity against law enforcement mandates, requires the kind of rigorous ethical analysis of competing interests provided in the following section.

Ethical Frameworks

The democratic harms of surveillance outlined above underscore why ethical analysis must move beyond privacy alone. If surveillance reshapes the conditions for collective action and civic trust, then evaluating its legitimacy requires a framework that weighs not only individual rights but also the structural integrity of democratic participation. This section examines the "security paradox" through utilitarian, moral-foundational, and deontological ethical lenses to identify underlying tensions and understand how it requires societies to balance anonymity and accountability within the broader pursuit of justice and public safety.

The Greater Good vs. Individual Rights

In the weeks and months following the attacks on September 11, 2001, many in the U.S. struggled to understand how such a monumental failure of its intelligence and law enforcement agencies could have occurred. Congress scrambled to introduce and overwhelmingly pass the Patriot Act, legislation that has shaped our collective approach to national security over the past twenty-five years (Goitein, 2021; Zuboff, 2019, pp. 113-115). In this new paradigm driven by fear, the prevention of terrorism outweighed the inconvenience of long airport check-in lines, tighter immigration control, "sneak peek" searches, and electronic surveillance. Advocates for the sweeping law enforcement legislation persuaded the public that restrictions on civil liberties

were necessary (Goitein, 2021; Zuboff, 2019, pp. 113-115). Exposure to these restrictions increased tolerance for state surveillance, normalized institutional power, and spawned programs such as DARPA's "Total Information Awareness," which would later become today's IVSS infrastructure (Harris, 2012).

Government agencies and private companies often defend IVSS through utilitarian reasoning, which seeks to maximize the greatest good for the greatest number. They argue that surveillance benefits society by preventing terrorism, solving crimes, and finding missing children, and that these benefits outweigh harms to privacy and autonomy. Bentham's "hedonic calculus" weighs pleasure and happiness against pain and suffering to determine the net social outcome (Hampton, 2023; Tavani, 2015, p. 45). In this model, "The moral value of actions and policies ought to be measured in terms of their social utility rather than via abstract criteria such as individual rights or social justice" (Tavani, 2015, p. 45). Bentham's "Panopticon" similarly shows how the threat of constant observation can shape behavior (Tavani, 2015, pp. 323–324). However, this reasoning overlooks the disproportionate burden surveillance places on vulnerable groups and the chilling effect it can have on democratic participation. As shown earlier, IVSS can discourage protest through movement tracking, facial recognition, data profiling, and device monitoring. Sacrificing privacy and constitutional rights for national security reduces citizens' agency and autonomy to means rather than ends (Tavani, 2015). Why, then, is public opinion on IVSS so divided?

The Moral Foundations of Public Attitudes Toward IVSS

Psychologist Jonathan Haidt's Moral Foundations Theory helps explain this divide. While utilitarian justifications for IVSS are rooted in abstract calculations about the greatest good, they often function as rationalizations for pre-existing intuitions. Moral Foundations Theory provides

a psychological framework for understanding how intuitions drive snap moral judgments. Haidt (2012) identifies six foundations that shape humans' ability to cooperate and make moral judgments, with gut intuition coming first and logical justification following. He defines each moral foundation through oppositional pairings that frame our interactions with one another and our understanding of our place within a just society. These foundations include Care, Fairness, Loyalty, Authority, Sanctity, and Liberty (Haidt, 2012).

Figure 5

Haidt's Six Moral Foundations

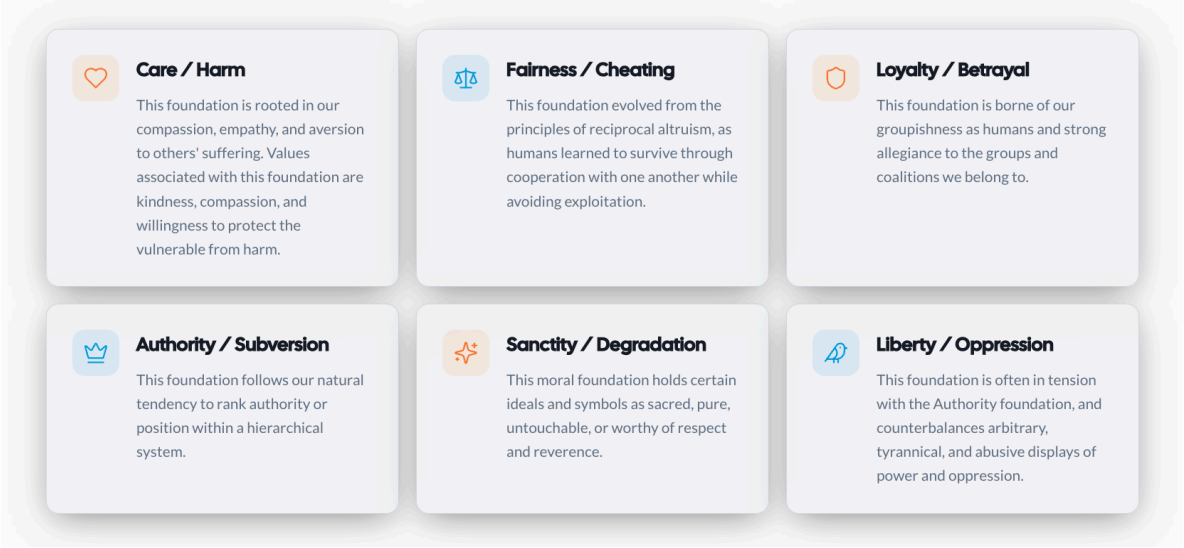

Members of the public often hold different moral and ethical views about the use of IVSS, how it is applied, and who becomes the target of scrutiny. People differ in the weight they assign to each moral foundation based on their worldviews, and these preferences shape their attitudes toward surveillance. Liberal and conservative views on IVSS often diverge according to how each side of the political spectrum weighs those foundations. According to Haidt, people with liberal political and social views tend to emphasize the individual-focused foundations of Care,

Fairness, and Liberty, while conservatives tend to emphasize the group-focused foundations of Loyalty, Authority, and Sanctity (Haidt, 2012).

The authors (Grandes et al., 2026) conducted a survey to investigate how people's moral foundations align with their views on intelligent surveillance systems. A total of 77 respondents answered basic demographic questions, including age, gender, location (urban, suburban, rural), country, level of education, and familiarity with IVSS. The respondents were asked to what degree they either agreed or disagreed with 12 statements about intelligent surveillance, two per Haidt foundation, one positive and one negative. Higher normalized scores for each foundation indicate strong anti-surveillance sentiments, while lower scores suggest pro-surveillance attitudes.

Finally, a series of questions asked about which concerns them more, governmental or corporate use of IVSS, and whom they trust more to self-regulate. Several large language models (LLMs) were used to identify and verify three meaningful segments, or *personas*, based on similar responses, using k-means clustering of the data. Analysis of the data identified three distinct personas according to their prevailing attitudes toward IVSS: Security-Forward Supporters, Conditional Pragmatists, and Civil Liberties Defenders (Grandes et al., 2026).

Figure 6

Average foundational scores per persona

Foundation	Civil Liberties Defenders	Conditional Pragmatists	Security-Forward Supporters
Care	4.35	3.32	2.53
Fairness	4.47	3.44	2.59

Loyalty	4.38	3.03	2.19
Authority	4.57	3.60	3.03
Sanctity	4.78	3.81	2.56
Liberty	4.53	3.55	2.41

Security-Forward Supporters

The Security-Forward Supporters (n = 16) persona is based on a cluster in the data with lower overall average scores, indicating that the segment is generally pro-surveillance and holds utilitarian attitudes toward IVSS. The Security-Forward Supporters' highest-scoring question underscored their belief that “Facial recognition technology is acceptable if it is applied equally to everyone”; however, they tend to overlook evidence that IVSS is likely to target certain groups unfairly. Both questions relate to the Fairness Foundation, which, in this case, could be interpreted to mean that people have nothing to fear unless they are cheating the system. Supporters also strongly believed that “Communities should support the use of facial recognition if it helps protect their members.” This statement reflects the values associated with the Loyalty Foundation, which prioritizes the community’s well-being over individual rights. Loyalty is a significant driver of inter-group competition, so taken in this context, the statement could be viewed as specifically protecting *their* community. High Sanctity and Authority Foundation scores for this cohort indicate that Security-Forward Supporters reject framing this topic as degrading or violating human dignity, the use of IVSS is perfectly acceptable, and law enforcement should have discretion in how it is applied. The only exception, based on comments from respondents in this segment, appears to be the use of IVSS by private corporations

engaging in the “widespread use of personalized advertisements” or “selling the information for profit” (Grandes et al., 2026).

Civil Liberties Defenders

Civil Liberty Defenders’ (n = 20) scores indicated vehement opposition to the use of IVSS technology to track and monitor citizens. Respondents in this group overwhelmingly agreed that facial recognition is likely to target certain groups and violate civil liberties. Answers reflect overall high indices in the Care, Liberty, and Fairness foundations. One respondent noted that they are “Concerned it will hurt innocent people,” and that IVSS is “not always accurate, especially with minorities, and people will be considered guilty without proof (Grandes et al., 2026). There was striking consistency across all moral foundations, with this cohort's scores reflecting views that this technology degrades human dignity (Sanctity Foundation) and that it betrays public trust (Loyalty and Fairness Foundations). There was unanimous support for governmental restrictions and for the view that law enforcement should not have discretion to apply IVSS in the field (Grandes et al., 2026).

Conditional Pragmatists

Conditional Pragmatists (n = 31) made up the largest of the three groups, with their scores sitting somewhere in the middle and generally reflecting the overall average across all participants. This cohort understood the technology's potential benefits and was willing to accept its use in certain contexts, but not without restrictions. Most scores indicated that they did not hold strong beliefs either way, with a few exceptions. One respondent expressed this general cautious acceptance, saying, “If we can track down children or suspects faster in public, then it may be worth it” (Grandes et al., 2026). However, there is a moderate level of mistrust in the

ability of both private and public institutions to implement the technology without causing harm. In fact, they remained highly skeptical, as indicated by beliefs that IVSS just feels wrong and that strong restrictions should be imposed on its use (Grandes et al., 2026).

Overall Takeaways

The study skewed heavily toward highly educated individuals living in suburban or urban settings, with almost all respondents (n=67) having earned a college degree. Most participants had at least baseline awareness of data collection and FRT, with over two-thirds reporting being moderately, very, or extremely familiar. Due to study limitations, it was not clear to what extent they *actually* knew about IVSS practices versus what they *thought* they knew. However, based on long-form responses, many had a sophisticated understanding of how the technology is used, its ingrained racial biases, lack of regulation, misuse, and mishandling of personally identifiable information, and, as one respondent posited, "disproportionately negatively affects BIPOC, elderly, and queer communities" (Grandes et al., 2026).

One finding stood out across all responses: irrespective of defined segments, respondents overwhelmingly showed skepticism about government overreach and corporate handling of user data. It was also clear that most people, regardless of persona groupings or demographics, mistrust the government and private entities to govern themselves effectively. Even those expressing support for IVSS added strong caveats for strict oversight and limitations of its use (Grandes et al., 2026).

Moral Foundation Theory provides insights into why people arrive at ethical justifications for or against this technology. However, it does not offer any moral clarity on whether IVSS can be deployed fairly in a just society. Utilitarianism offers a framework to calculate the net benefit or detriment of IVSS to society, but it fails to justify collateral damage to

specific vulnerable populations and the chilling effect on civic participation. Applying a deontological lens to the security vs. privacy paradox can help shed light on why privacy is a fundamental human right.

Privacy and Autonomy as Human Rights

Privacy is not a simple concept to define because it affects so many aspects of daily life, yet it is a fundamental human right that is core to a thriving, democratic society (Roessler & DeCew, 2023). Philosopher John Rawls (1999, p. 60) argues that in a just society, everyone has an equal claim to "equal basic liberties" such as the freedom of speech, association, and privacy. Rawls (1999, p. 136) proposes a thought experiment, imagining a "veil of ignorance," in which one is unaware of their status or position in society. Under this veil, would any rational person agree to participate in a surveillance system more likely to misidentify them by skin tone? Would anyone agree to be permanently profiled in a law enforcement database? The likely answer is that they would not.

Privacy is non-negotiable for the functioning of liberal societies because, without it, citizens participate at their own risk. According to Ruth Gavinson (1980), Privacy is crucial for the functioning of liberal societies because it “fosters and encourages the moral autonomy of the citizen, a central requirement of a democracy.” Conversely, the “unequal distribution of privacy may lead to manipulation, deception, and threats to autonomy and democracy” (Gavison, 1980, p. 421). When asked what concerns them most about the use of IVSS technology, one anonymous survey respondent commented on “These resources being used or exploited for greed and/or control… potentially destroying the freedoms of individuals (Grandes et al., 2026)” When a society experiences the threat of constant surveillance, its citizens are subject to manipulation and coercion, as were the prisoners in Bentham’s Panopticon (Tavani, 2015, p. 324). Shoshana

Zuboff (2019, p. 469) argues that this exploitation and commoditization of human nature, "reduces us to our behavior, transformed into another fictional commodity and packaged for others' consumption." That is precisely the threat as surveillance infrastructure creates asymmetrical power dynamics between those who collect and monitor data, and those subject to collection. She argues that the power wielded by the surveillance state is no longer about controlling data about us, but rather about using it to shape our behavior (Zuboff, 2019, p. 189).

Shared Responsibility and the Need for Governance

Benjamin Franklin famously declared, "They who can give up essential liberty to obtain a little temporary safety, deserve neither liberty nor safety" (Franklin's Contributions to the Conference on February 17: Four Drafts, 1775, 1978). The view that privacy is an inalienable human right may seem irreconcilable with a utilitarian approach to collective security, but the security paradox is not really a paradox at all; it is a false dilemma. While this framing captures the core tension at the heart of the security paradox, it assumes that a surplus in security is directly proportional to a deficit in liberty, and vice versa. Kirchschläger (2019) insists that when two rights conflict, neither can be chosen over the other because both are essential to human existence. Instead, they should be thought of as indivisible, with humans as the bearers of *all* human rights (Kirchschläger, 2021). Furthermore, Nissenbaum (2009) suggests that privacy is not about keeping information secret, but about the responsible flow of information. Violations only occur when information is disseminated without appropriate permissions, for inappropriate purposes, or in violation of applicable laws and norms (Nissenbaum, 2009).

As power and influence become concentrated in the hands of a shrinking number of actors, it is increasingly important to share responsibility for reigning in harms and maximizing benefits (Kirchschläger, 2021). The government bears the greatest responsibility in ensuring

constitutional rights are protected through legislation and enforcement. Private companies must be held accountable for the ethical design of these systems and the responsible handling of sensitive data. For example, privacy advocacy organizations like the Electronic Privacy Information Center, also known as EPIC (2026), and the United Nations Educational, Scientific, and Cultural Organization, or UNESCO (2022), have proposed that governance of privacy is possible through meaningful oversight of data collection and processing, and through explainability and transparency. Finally, citizens play a critical role by staying informed, advocating for their rights, and resisting the normalization of privacy infringement by exercising control over their data and engaging in democratic participation to pressure power structures.

Public trust in big tech companies and the government to self-govern remains low and insufficient to prevent abuses of power. Without incentives for compliance, enforcement of restrictions, and real accountability, it is not readily apparent how an ethical and moral balance between security and civil liberties can be achieved. Still, there is room to preserve both civil liberties and security through the combined efforts of all responsible parties, but doing so requires appropriate governance frameworks and legislation to ensure that the balance is fair.

Translating Ethical Principles into Enforceable Obligations

Ethical principles are not fixed foundations for regulating IVSS; they function as high‑level abstractions that require contextual interpretation and negotiation among competing interests (Floridi & Cowls, 2019; Bleher & Braun, 2023; Munn, 2023). The legal rules that emerge from this process reflect political compromise and the relative power of stakeholders within each system (Nelson, 2026; Justo‑Hanani, 2026). Western nations have similar ethical frameworks and cultural values. Examining the difference in regulation of AI between the United

States and Europe demonstrates the importance of political structure, regulatory philosophy, and the role of the state in mediating technological risk.

The EU's regulatory model—rooted in fundamental rights, precaution, and strong public-interest institutions—treats biometric and behavioral surveillance as inherently high-risk and subject to strict ex ante controls. By contrast, the United States relies more heavily on market forces, sector-specific rules, and after-the-fact remedies, producing a fragmented and permissive environment for IVSS deployment. The following sections compare these approaches to show how ethical principles are translated into enforceable obligations within each political economy.

Political Economy of AI Regulation

Understanding why the EU and the United States regulate IVSS so differently requires examining the political-economic forces shaping how each system converts ethical principles into law. The United States is home to many of the world's dominant technology companies, and these firms have long exerted substantial influence over federal policy. The rise of AI, which requires massive capital investment and therefore the promise of equally massive returns, has intensified incentives for political engagement. Many firms strategically shifted their political affiliations to maximize leverage (Nelson, 2026), while the removal of limits on private election spending enabled them to deploy virtually unlimited resources to shape regulatory outcomes (Justo-Hanani, 2026).

In the IVSS context, this produced a new form of public–private alignment. Technology companies develop systems capable of extracting and analyzing personal information at scale, effectively generating a dossier on each person. Government agencies adopt these systems not only for law enforcement but also to monitor and deter dissent, as seen in the surveillance of

ICE-related protests. This is not traditional regulatory capture but a mutually reinforcing partnership: the state gains unprecedented visibility into the population, while firms gain political protection, market dominance, and regulatory indulgence. As Zuboff argues, surveillance capitalism "produces new markets of behavioral prediction and modification," and the government becomes a central beneficiary. Such systems can enhance networked authoritarianism by manipulating democratic opinion-formation and weakening pluralism (Kirchschlaeger, 2021).

This alignment operates through two reinforcing flows of power. First, IVSS enables the government to learn extraordinary amounts about individuals, their movements, associations, habits, and networks. Second, the same infrastructure allows the government to shape the information environment itself. AI-driven personalization gives platforms the ability to amplify preferred narratives and mute criticism without producing content directly. Personalization may manipulate how humans use their right to seek information and their right to form an opinion, undermining democratic deliberation (Kirchschaler, 2021). The scale, speed, and opacity of these systems widen the "governance gap" between technological capability and legal constraint.

By contrast, the EU's more diffuse technology sector and stronger public-interest regulatory institutions make precautionary regulation politically feasible. The EU's approach, exemplified by the GDPR and the AI Act, treats biometric surveillance as a high-risk activity requiring strict purpose limitation, data minimization, and ex ante authorization. These structural differences explain why the United States lacks comprehensive federal regulation, why federal proposals include broad preemption of state privacy laws, and why IVSS has expanded rapidly despite its risks to civil liberties and democratic participation.

The EU AI Act

The nations that comprise the EU, for the most part, are liberal democracies that value autonomy and democratic participation. The EU identified and sought to address harms from computer surveillance. The EU AI Act is the most ambitious attempt to convert ethical principles into binding obligations for high-risk technologies such as IVSS. It adopts a risk-based governance framework that combines targeted prohibitions with extensive procedural safeguards. The Act classifies AI systems by the risks they pose to health, safety, and fundamental rights (Regulation (EU) 2024/1689). IVSS is designated "high-risk" across public services, law enforcement, and border control (Annex III).

The Act treats mass biometric surveillance as an unacceptable risk. Recital 32 warns that emerging technologies can create a "feeling of constant surveillance," and Article 5 bans real-time remote biometric identification in public spaces, with narrow exceptions such as locating a missing child or preventing a terrorist attack. Critics argue the Act focuses too narrowly on biometric identification and does not fully address behavioral analytics that track individuals without identifying them (Veale & Zuiderveen Borgesius, 2021).

Article 5 prohibits untargeted scraping of facial images from CCTV or the internet, directly targeting Clearview-style models, and bans AI systems that classify individuals based on social behavior or personality traits. Real-time biometric identification is allowed only for serious, narrowly defined purposes and requires prior judicial authorization.

The Act requires high-risk systems to undergo pre-deployment testing to identify and mitigate bias (Arts. 10, 15), reflecting research showing demographic disparities in facial-recognition accuracy (Buolamwini & Gebru, 2018). Deployers must assess impacts on

privacy and equality (Art. 27), and GDPR reinforces strict purpose-limitation and data-minimization requirements. However, GDPR remains a data-processing framework, not a comprehensive response to behavioral prediction or population-level inference (Cohen, 2013; Zuboff, 2015). To address transparency concerns, the Act mandates documentation, human oversight, and mechanisms for individuals to contest automated decisions (Arts. 13–15; Laux et al., 2023).

United States Approaches

U.S. AI policy has shifted from emphasizing human rights and safety in the Biden Administration to the Trump Administration's belief that any regulation risks impeding innovation. As a result, current limits on IVSS arise primarily from constitutional doctrine and a patchwork of state and municipal laws. In March 2026, competing federal proposals were introduced in Congress, but given political polarization, it is unclear whether any legislation will pass during the second Trump administration.

Federal Regulatory Landscape

The Biden administration's *Blueprint for an AI Bill of Rights* and Executive Order 14110 (2023) emphasized rights-protective principles: safe and effective systems, prevention of algorithmic discrimination, data-privacy safeguards, and human oversight. The Trump administration's *America's AI Action Plan* (2025) reframed the federal role as a facilitator of technological advancement and national security, emphasizing deregulation, accelerated deployment, especially in border enforcement, and removal of equity requirements (Executive Office of the President, Office of Science and Technology Policy, 2025; Nelson, 2026). While the shift in approach was vast in many ways, it did not end the debate on the role of regulation.

Proponents of regulation argue that more regulation will enhance trust in AI and thereby support innovation. Each side has significant public support and, as noted below, the debate over the role of regulation continues at both the federal and state levels.

President Trump's approach was codified in the *One Big Beautiful Bill Act* (OBBBA), enacted July 4, 2025 (Congress.gov, 2025). OBBBA allocated billions for IVSS infrastructure, including Autonomous Surveillance Towers and AI-driven border scanning (Congressional Budget Office, 2025). Although the bill originally included a 10-year moratorium preempting all state AI regulations, the Senate removed this provision (U.S. Senate, 2025). The administration then issued Executive Order 14365, establishing an AI Litigation Task Force to challenge restrictive state laws and authorizing the withholding of federal broadband funds from states with stringent privacy standards (Exec. Order No. 14,365, 2025). The executive order reflects Trump's first principle that there be no regulation of AI. A federal task force is set up to challenge state regulation. The withholding of federal funds to influence state regulation is also a tool frequently used by the Trump administration.

Proposed Federal Legislation Affecting IVSS

On March 20, 2026, the White House released its *National AI Legislative Framework*, arguing that a "patchwork of conflicting state laws" should be replaced with uniform national rules (White House, 2026). The framework expands federal authority to combat AI-enabled scams, strengthens national-security uses of AI, and streamlines permitting for data-center policies that indirectly facilitate large-scale surveillance.

Senator Blackburn's RUMPAI Act adopts this deregulatory posture. Title XVII includes a broad preemption clause overriding state biometric-privacy statutes, which are currently the

strongest limits on IVSS data collection and retention (Blackburn, 2026). Title VII creates a federal cause of action for consumer-facing AI harms but excludes government surveillance technologies, leaving IVSS-related harms outside the bill's remedial structure.

In contrast, Senator Wyden's *Government Surveillance Reform Act of 2026* would impose sweeping constraints on federal surveillance. It requires agencies to obtain a judicial warrant before accessing sensitive data, including location information, web-browsing and search records, chatbot logs, car telematics data, and comparable social-media-derived data, whether acquired directly or purchased from brokers, closing the "data-broker loophole" (Wyden, 2026).

Constitutional Constraints on Use of IVSS by Law Enforcement

The Constitution constrains government use of IVSS, but existing doctrine is poorly equipped to regulate pervasive, automated surveillance. Fourth Amendment jurisprudence, which is grounded in privacy and protection from arbitrary intrusion, must now be applied to technologies unimaginable to the Framers (Slobogin & Brayne, 2023). Under the "reasonable expectation of privacy" test, courts have long held that individuals lack privacy in what they expose to the public or share with third parties (Katz, 1967), with only narrow exceptions such as cell-site location information (Carpenter, 2018).

These doctrines allow police to use CCTV, facial recognition, and services like Clearview AI without triggering a "search," and to obtain footage from private smart-home devices without a warrant (Wang et al., 2024). Scholars warn that IVSS transforms public-space observation into continuous, retrospective tracking that approaches the kind of general search the Fourth Amendment forbids (Ferguson, 2025; Kerr, 2018; Tuggle, 2021). IVSS functions as a "visual superpower" and a "time machine," enabling retrospective analysis of everyone, everywhere, all

at once (Ferguson, 2025). Given the Supreme Court's deference to law enforcement, a categorical ban is unlikely.

The Fourth Amendment doctrine alone cannot capture the constitutional stakes. IVSS burdens the freedoms of speech, assembly, and association by enabling identity-linked tracking of individuals engaged in expressive activity. Public disclosure of protest surveillance "renders individual protestors identifiable" and exposes them to "threats of private retaliation," chilling participation (Valeska, 2021). In *Americans for Prosperity Foundation v. Bonta*, the Court held that compelled disclosure of associational information imposes a cognizable First Amendment injury because it creates "an unnecessary risk of chilling" protected expression. IVSS-enabled protest surveillance is far more intrusive than the confidential donor-reporting regime struck down in *Bonta*.

The doctrinal gap is stark. *Laird v. Tatum* (1972) insulated passive information collection from First Amendment challenge, but that case involved speculative harms and did not address public identification or dissemination. IVSS is categorically different: it enables the government to identify protestors, track their movements, and publicize their participation. Federal agencies, including ICE, have used computer-vision tools, social-media scraping, and data-broker purchases to monitor political gatherings and immigrant-rights protests. Such practices implicate the core of First Amendment protection: the right to engage in collective political expression without fear of surveillance or reprisal.

IVSS thus represents a new paradigm of state power that the existing Fourth Amendment doctrine does not meaningfully address because the harm is not merely privacy invasion and misidentification by law enforcement, but suppression of political participation. Any

constitutional analysis must integrate First Amendment principles alongside Fourth Amendment doctrine.

Patchwork of State Regulation

Because neither federal regulation nor constitutional doctrine meaningfully restricts IVSS, states and municipalities have stepped in to fill the gap. State regulation operates primarily through biometric-privacy statutes and broader data-privacy laws. The most influential is the Illinois Biometric Information Privacy Act [BIPA] (2024), which prohibits the sale of biometric data and requires notice, consent, and public retention schedules. Texas and Washington have enacted similar but less stringent statutes.

These protections are now at risk. Blackburn (2026) contains a sweeping federal preemption clause that would override state biometric-privacy statutes, including BIPA. If enacted, these provisions would eliminate the only meaningful constraints on IVSS, shifting governance to federal agencies that have historically prioritized innovation and national security over privacy. BIPA's private right of action has made it the most consequential state-level constraint on IVSS. The ACLU's lawsuit against Clearview AI proceeded solely under BIPA and resulted in a settlement imposing enforceable limits on Clearview's scraping and use of biometric data. Federal preemption would eliminate this enforcement mechanism entirely (American Civil Liberties Union v. Clearview AI, Inc., 2022).

Proposed Comprehensive Legislation to Govern IVSS

The rapid expansion of IVSS has outpaced the legal frameworks designed to protect privacy, civil liberties, and democratic participation. Existing doctrine does not adequately

constrain systems capable of persistent monitoring, biometric identification, retrospective search, and cross-platform data fusion (Garvie et al., 2016). Drawing on proposals to protect demonstrators from government surveillance (ICNL, 2023) and privacy-first regulatory models (EFF, 2023), this section outlines a legislative framework for regulating IVSS across law-enforcement and non-law-enforcement contexts.

Core Principles

A rights-protective IVSS statute should be grounded in four principles:
(1) an expanded right to privacy recognizing IVSS's capacity to generate lifelong dossiers;
(2) strict limits on government surveillance, especially during expressive activity (ICNL, 2023);
(3) data minimization and purpose limitation (GDPR, 2016, art. 5); and
(4) transparency, accountability, and enforceability (Garvie et al., 2016).

Applying these principles, Congress should deem that using IVSS to track a specific individual's movements, including accessing historical IVSS data, facial-recognition matches, license-plate reader logs, or geolocation-derived patterns, constitutes a Fourth Amendment search. A judicial warrant based on probable cause must be required before law enforcement may use IVSS to identify, track, or locate an individual. Probable cause for an arrest warrant may not be based solely on IVSS data; any match must be corroborated by independent evidence (ICNL, 2023).

Congress should prohibit the use of IVSS to identify, track, or catalog individuals engaged in lawful First Amendment activity absent individualized, fact-based suspicion of criminal wrongdoing (ICNL, 2023). Any such use must receive supervisory approval and be strictly time-limited (Franceschi-Bicchierai & Whittaker, 2026; Frenkel & Krolik, 2026).

Moreover, collection, retention, and use of biometric and behavioral data must be limited to what is strictly necessary for an authorized purpose, defined as a specified, explicit, and legitimate purpose within the agency's statutory authority, and may not be further processed in a manner incompatible with that authority (GDPR, 2016, art 5). Data must be deleted once the purpose is fulfilled. Agencies may not purchase biometric or geolocation data from brokers (Electronic Frontier Foundation, 2023). Private entities may transfer such data only with informed, specific, opt-in consent. Individuals must have the right to access, correct, delete, and port their data.

Agencies must also publicly disclose all IVSS tools in use and the purpose of each deployment. Vendor contracts must be made public within 30 days. Annual reports must document accuracy problems, unauthorized uses, and results of independent audits, including demographic impacts (Garvie et al., 2016).

The statute should include a private right of action allowing individuals to seek statutory and compensatory damages, injunctive relief, and attorneys' fees (EFF, 2023).

Finally, Federal legislation should not preempt stronger state protections unless the federal statute is more comprehensive. Pending federal proposals include broad preemption provisions that would weaken existing safeguards; such provisions should be removed unless Congress enacts a more protective framework.

Conclusion

This paper has traced the security paradox that IVSS purveys by burdening the very citizens it seeks to protect, from the technical architecture of Intelligent Vision Surveillance Systems, through their documented harms of privacy erosion, false identification, discriminatory

burden, and cybersecurity vulnerability, to the ethical frameworks that expose these harms as ethical quandaries, and finally to the legal obligations those frameworks demand. Taken together, these sections establish that IVSS is a system whose structural design may produce injustices, alter democratic expression, and one that existing law has not yet been adequately reformed to address.

The emerging frontier of mass AI surveillance makes these conclusions more urgent, not less. In July 2025, Anthropic signed a two-year, $200 million contract with the U.S. Department of Defense to prototype frontier AI capabilities for national security (NBC News, 2026; Anthropic, 2025). The partnership collapsed publicly in February 2026, when the Pentagon demanded that Anthropic remove contractual safeguards prohibiting the use of its Claude models for mass domestic surveillance of American citizens and fully autonomous weapons systems. Anthropic refused, and the Trump administration responded by ordering all federal agencies to immediately cease use of Anthropic's technology and designating the company a national security supply chain risk (NPR, 2026a; NPR, 2026b). The ongoing dispute lays bare what this paper has argued from the outset: the pressure to deploy AI at scale for security purposes consistently conflicts with the rights-protective limits that ethical and legal analysis demands.

IVSS, as analyzed here, represents the current frontier of this tension. Mass AI surveillance, integrating large language models, predictive behavioral analysis, and real-time biometric identification across entire populations, represents the next. The novel harms that frontier may introduce, including AI-generated threat assessments that target dissent at scale, fully automated detention decisions, and the erosion of privacy, urgently require scholarly attention. The work of understanding, challenging, and governing what comes next has only begun.

As discussed, there is currently no comprehensive AI regulation in the United States, though Congress will continue to consider options. Advancing meaningful reform will require research grounded in robust datasets that can assess the cumulative democratic effects of IVSS deployment—tracking measurable impacts on participation, protest activity, and political inequality. Yet evidence alone will not drive legislative change. Limiting IVSS use demands institutional reform and sustained public pressure for participatory rulemaking that gives civil liberties genuine weight in policy design. Scholars can strengthen this effort by addressing the security paradox: how to preserve fundamental rights while deploying novel technologies to enhance collective security. The paradox does not require an all-or-nothing solution but for a principled equilibrium that treats both privacy and safety pillars of democratic governance.